arXiv:1612.07272v1 [physics.data-an] 21 Dec 2016

Probabilistic prediction of the AL index with the diffusion forecasting model

Dimitrios Giannakis,[1] Matina Gkioulidou,[2] and John Harlim[3]

[1] *Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA*
[2] *The Johns Hopkins University, Applied Physics Laboratory, Laurel MD 20723-6099, USA*
[3] *Department of Mathematics and Department of Meteorology and Atmospheric Science, The Pennsylvania State University, University Park, PA 16802-6400, USA*

(Dated: December 22, 2016)

We propose a nonparametric approach for probabilistic prediction of the AL index trained with AL and solar wind (vB_z) data. Our framework relies on the diffusion forecasting technique, which views AL and vB_z data as observables of an autonomous, ergodic, stochastic dynamical system operating on a manifold. Diffusion forecasting builds a data-driven representation of the Markov semigroup governing the evolution of probability measures of the dynamical system. In particular, the Markov semigroup operator is represented in an orthonormal basis acquired from data using the diffusion maps algorithm and Takens delay embeddings. This representation of the evolution semigroup is used in conjunction with a Bayesian filtering algorithm for forecast initialization to predict the probability that the AL index is less than a user-selected threshold over arbitrary lead times and without requiring exogenous inputs. We find that the model produces skillful forecasts out to at least two-hour leads despite gaps in the training data.

The interaction of the solar wind (SW) and the embedded interplanetary magnetic field (IMF) with the intrinsic dipolar geomagnetic field creates a cavity in the Earth's magnetosphere. Southward turning of the IMF allows for the efficient penetration and storage of SW energy in that cavity, whose subsequent release leads to an explosive phenomenon known as magnetospheric substorm. During substorms, a strong ionospheric current is developed, inducing geomagnetic field disturbances on the ground. The intensity of a substorm is quantified by the AL index, which is a measure of the magnitude of those disturbances. Understanding and predicting geomagnetic substorms, and thus one of the major modes of SW–magnetosphere coupling, has been a long-standing, unresolved issue in space physics. This is because substorms involve multiscale, nonlinear interactions of plasmas that are either not in equilibrium, or are in unstable equilibrium. Largely due to the multiscale nature of the phenomenon, first-principles physics-based modeling is not always feasible and often produces erroneous predictions of the magnetospheric state during substorms. Therefore, developing data-driven predictive models that take advantage of the abundance of AL index measurements is of the essence. Various data-driven techniques have been proposed, such as non-autonomous models [1] and techniques based on support vector machines [2]. Frequently, SW conditions are used as external factors in order to predict the AL index. In particular, the external input is the product vB_z where v is the bulk SW velocity along the Sun-Earth axis and B_z the north-south component of the IMF. As stated above, when the IMF is southward ($B_z < 0$), the SW–magnetosphere coupling becomes very efficient and more likely to cause a geomagnetic substorm, and thus decrease of the AL index.

In this paper, we propose a nonparametric approach for statistical prediction of AL index time series using a recently developed data-driven technique called diffusion forecasting [3]. This approach models the AL index as an observable (function of the state) of an abstract stochastic dynamical system operating on a manifold. In diffusion forecasting, the semigroup operator governing the evolution of probability measures is represented on a smooth data-driven basis acquired from time-ordered data using the diffusion maps algorithm [4] with a variable-bandwidth kernel [5]. Using this representation of the dynamics and a Bayesian filtering scheme for forecast initialization [6], our method predicts the probability for the observable of interest (in this case, the AL index) to lie in given intervals for arbitrary lead times without requiring exogenous inputs. In particular, unlike other approaches [1], our model does not treat vB_z as an external factor for AL index prediction but rather as an observable of a single autonomous dynamical system. Thus, as an auxiliary result, our model also predicts analogous probabilities for the vB_z data (see supporting information (SI)).

Consider time-ordered datasets, $\{z_0^{\rm AL}, z_1^{\rm AL}, \ldots, z_{N-1}^{\rm AL}\}$ and $\{z_0^{\rm SW}, z_1^{\rm SW}, \ldots, z_{N-1}^{\rm SW}\}$, of AL index and vB_z time series, respectively, taken at times $\{t_i = i\,\tau + t_{\rm ref}\}_{i=0}^{N-1}$ for a fixed sampling interval τ and starting time $t_{\rm ref}$, and normalized to unit variance. In our training procedure we use $\tau = 5$ min and $N = 20{,}000$ corresponding to a physical interval of approximately 69 days. Here, we focus on results from a training interval starting on 1/1/2003; a 10-day portion of the AL index and vB_z data in that interval is shown in Fig. 1. Our data sources are listed in the SI. Note that the vB_z data sometimes contain gaps, which we fill in using linear interpolation in time. As demonstrated below, diffusion forecasting is able to produce skillful forecasts despite such missing data. Results from a training interval starting on 1/1/2000 are presented in the SI.

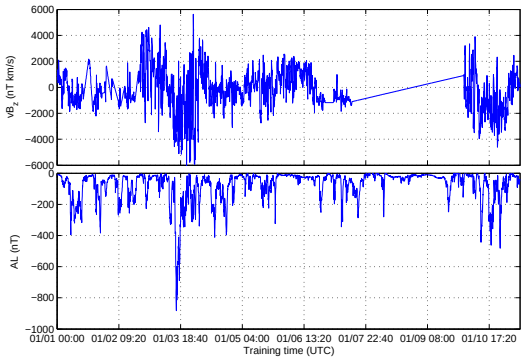

FIG. 1. vB_z (top) and AL index (bottom) time series for the interval 1/1/2003 0:00 UTC to 1/11/2003 10:00 UTC. Notice the missing SW data occurring around the 1/07 22:40 UTC tickmark which are filled in using linear interpolation in time.

The diffusion forecasting approach views the AL index and vB_z time series as observables on the state space manifold M of an ergodic (hence autonomous) stochastic dynamical system; that is, we consider that there exist smooth mappings $h^{\mathrm{AL}} : M \mapsto \mathbb{R}$ and $h^{\mathrm{SW}} : M \mapsto \mathbb{R}$ such that $z_i^{\mathrm{AL}} = h^{\mathrm{AL}}(x_i)$ and $z_i^{\mathrm{SW}} = h^{\mathrm{SW}}(x_i)$, where x_i is the system state at time t_i. Similarly, we have the composite map $h : M \mapsto \mathbb{R}^2$ with $h(x_i) = z_i = (z_i^{\mathrm{AL}}, z_i^{\mathrm{SW}})$. In general, knowledge of z_i alone is not sufficient to uniquely determine the underlying state x_i. Nevertheless, according to the theory of delay-coordinate maps [7], for a sufficiently large number of delays q and under mild assumptions on f, δt, and the dynamical system, the map $H : M \mapsto \mathbb{R}^{2q}$ with $H(x_i) = y_i = (z_i, z_{i-1}, \ldots, z_{i-q+1})$ is an embedding of M, i.e., a one-to-one map with continuous inverse on its image. Thus, the data $y_i \in \mathbb{R}^{2q}$ are in one-to-one correspondence with the states $x_i \in M$. In what follows, we set $q = 10$ to give the model the information in a 50-minute interval; this is the typical time delay between the observed solar wind southward turning and the AL index decreases. Another crucial property that we make use of is ergodicity—this allows us to approximate expectation values with respect to the invariant probability measure of the dynamics by time averages. In particular, ergodicity implies that for an integrable observable f, $N^{-1}\sum_{i=0}^{N-1} f(x_i) \xrightarrow{\text{a.s.}} \int_M f\, d\mu$, where μ is an invariant probability measure.

Suppose now that at forecast initialization time knowledge about the system state is characterized by a probability measure with density ρ_0 relative to μ, and that ρ_0 lies in the Hilbert space $L^2(M,\mu)$. At lead time τ, the probability density for the system state x is given by $\rho_\tau(x) = e^{\tau\mathcal{L}^*}\rho_0(x)$, where $\mathcal{L}^*$ is the Fokker-Planck operator governing the evolution of probability densities with respect to μ. The key observation made in [3] (which is a consequence of ergodicity) is that given two functions $f_1, f_2 \in L^2(M,\mu)$ whose values $f_j(x_i)$ at the sampled states are known, the time-shifted correlation $N^{-1}\sum_{j=0}^{N-1} f_1(x_i) f_2(x_{i+1})$ is an unbiased estimate of the inner product $\langle f_1, e^{\tau\mathcal{L}} f_2\rangle = \int_M f_1 e^{\tau\mathcal{L}} f_2\, d\mu$ on $L^2(M,\mu)$ with variance scaling as N^{-1}. Here, $\mathcal{L}$ is the Kolmogorov generator of the dynamics, i.e., the adjoint of the Fokker-Planck operator $\mathcal{L}^*$. In particular, given a finite orthonormal set $\{\phi_0, \ldots, \phi_{\ell-1}\}$ in $L^2(M,\mu)$, we can compute the matrix elements $A_{ij} = N^{-1}\sum_{k=0}^{N-1} \phi_i(x_k)\phi_j(x_{k+1}) \approx \langle \phi_i, e^{\tau\mathcal{L}}\phi_j\rangle$. The key idea in diffusion forecasting is to approximate all densities by their projections onto the finite-dimensional subspace of $L^2(M,\mu)$ spanned by the ϕ_i; that is $\rho_\tau \approx \hat{\rho}_\tau = \sum_{i=0}^{\ell}\langle \phi_i, \rho_\tau\rangle\phi_i$. At lead time $t = m\tau$, $m \in \mathbb{N}$, we obtain the forecast density

$$\hat{\rho}_t(x_k) = \textstyle\sum_{i,j=0}^{\ell-1} c_j [A^m]_{ji}\phi_i(x_k), \qquad (1)$$

where $c_j = N^{-1}\sum_{k=0}^{N-1}\phi_j(x_k)\rho_0(x_k) \approx \langle \phi_j, \rho_0\rangle$ approximate the expansion coefficients at initial time. As is evident from (1), $\hat{\rho}_t(x_k)$ can be evaluated for arbitrary t without requiring exogenous inputs, so long as the expansion coefficients c_j are known. Moreover, no restrictions on the stepsize τ are placed.

In diffusion forecasting, the basis functions $\{\phi_i\}$ are computed from the dataset $\{y_i\}$ using the diffusion maps algorithm with a variable-bandwidth kernel designed to induce orthogonality of the ϕ_i with respect to the invariant measure. Moreover, the basis satisfies an optimality condition for the variance of the matrix elements A_{ij}. In summary, diffusion forecasting can be thought of as a spectral method for the dynamical evolution semigroup in a smooth data-driven basis inherited from the dataset. We refer the reader to [3] for further details on the construction of the basis and for error bounds. In our experiments, we set the number of basis functions to $\ell = 100$; the attained forecast skill is not too sensitive on this choice. We also build univariate diffusion forecast models using delay-coordinate mapped AL index data alone. We use such models to assess the contribution of the vB_z data in the prediction skill of our bivariate models based on joint AL index and vB_z observations.

As with any forecasting technique, the skill of diffusion forecasting depends strongly on the initial density ρ_0. Here, we perform forecast initialization using the Bayesian filtering approach developed in [6]. In particular, let $u_0 \in M$ be the (unknown) system state at forecast initialization time, $t = 0$, and $v_0 = h(u_0) \in \mathbb{R}^2$ be the corresponding AL index and SW data. Let also $\{u_{-n}, u_{-n+1}, \ldots, u_{-1}\}$ and $\{v_{-n}, v_{-n+1}, \ldots, v_{-1}\}$ with $v_i = h(u_i)$ be the system states at the n preceding timesteps to forecast initialization where n represents a filter spinup time. The scheme of [6] utilizes the familiar predictor-corrector approach from data assimilation to construct an initial density $\hat{\rho}_0(x_k)$ on the training data $\{x_k\}$ that approximates the conditional density $\rho(x_k \mid v_{-n}, \ldots, v_0)$. The method employs a Gaussian likelihood $L(v, x_k) = e^{-\|v-y_k\|^2/(2\sigma^2)}$, where the vari-

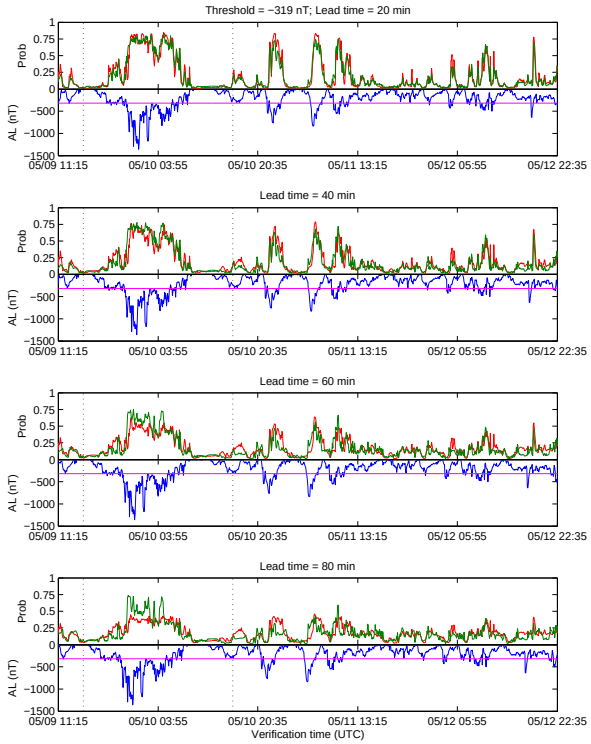

FIG. 2. Prediction of the AL time series for the interval 5/9 11:15 UTC to 5/12 22:35 UTC. Each plot shows the true time series (blue) and the prediction $p_t^{\mathrm{AL}}(a)$ for $a = -319$ nT (horizontal magenta line) at lead times $t = 20, 40, 60, 80$ minutes. Green and red lines show the predictions based on the bivariate and univariate models, respectively. The vertical doted lines indicate the interval shown in detail in Fig. 3.

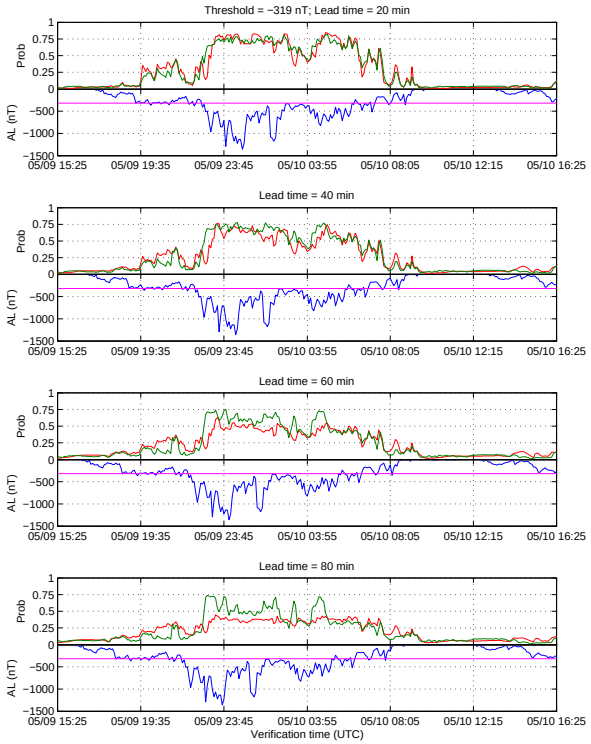

FIG. 3. Same as Fig. 2, but for the interval 5/9/2003 11:15 UTC to 5/12/2003 22:35 UTC. Note the skill improvement of the bivariate model trained on AL and vB_z data (green line) over the univariate model trained on AL data alone (red line).

ance σ^2 is estimated from the reconstruction error of the training data for the ℓ eigenfunctions used in the diffusion forecast model; i.e., $\sigma^2 \approx \int_M \|h - \hat{h}\|^2 \, d\mu$, where $\hat{h} = \sum_{i=0}^{\ell-1} \langle \phi_i, H \rangle \phi_i$. In the first filtering timestep we use a uniform prior, $\hat{\rho}_{-n}(x_k) = 1/N$, which we update to the posterior $\hat{\rho}^+_{-n}(x_k) = L(v_{-n}, x_k)\hat{\rho}_{-n}(x_k)/Z_{-n}$ using Bayes' rule (here, Z_{-n} is a normalization factor computed via ergodic averages). The density $\hat{\rho}^+_{-n}(x_k)$ is then updated using the diffusion forecast model to obtain the prior $\hat{\rho}_{-n+1}(x_k)$ via (1). Repeating this procedure $n-1$ times, we arrive at the density $\hat{\rho}^+_0(x_k)$ which we use as the initial condition for the diffusion forecast model. Given the predicted probability densities from diffusion forecasting, we perform probabilistic AL index prediction by computing the probabilities $p_t^{\mathrm{AL}}(a)$ for the AL index to be smaller than preselected thresholds a. These probabilities are given by $p_t^{\mathrm{AL}}(a) = \int_M \rho_t \chi_a \, d\mu$, where χ_a is the indicator function of the subset $\{x \in M : h^{\mathrm{AL}}(x) \le a\}$. Numerically, we approximate $p_t^{\mathrm{AL}}(a)$ by identifying the timestamps $T_a = \{j \in \{0, \ldots, N-1\} : h^{\mathrm{AL}}(x_j) \le a\}$ in the training dataset where the AL index is smaller than the given threshold, and evaluating $p_t^{\mathrm{AL}}(a) = N^{-1} \sum_{j \in T_a} \hat{\rho}_t(x_j)$.

Figure 2 displays the probabilities $p_t^{\mathrm{AL}}(a)$ for running forecasts from 3/1/2003 0:35 UTC to 3/4/2003 11:55 UTC (a disjoint interval from the training data) and lead times t in the range 20–80 minutes. In each panel, we show the AL index superimposed with the probabilities $p_t^{\mathrm{AL}}(a)$ for the univariate and bivariate models trained by the AL index alone and joint (AL, vB_z) data, respectively. The threshold is $a = -319$ nT which corresponds to one standard deviation below the AL index mean. In Fig. 2, significant decreases of the AL index have a clear correspondence with high probability, predicted by our model, that the AL magnitude will cross the threshold at lead t. Moreover, it is evident that the bivariate model provides improved long-term forecasts. For instance, for the substorm shown in detail in Fig. 3, the bivariate model gives predictions with higher probabilities compared to those of the univariate model when the AL index is actually smaller than a and lower probabilities when the AL exceeds a. These differences persist even for smaller thresholds (see SI). The higher skill of the bivariate model is to be expected since magnetospheric substorms are highly driven by changes in the solar wind.

To compare the performance of these two models in a more quantitative manner, we compute the receiver operating characteristic (ROC) used for evaluating a binary classifier system. Here, we use the ROC to assess the skill

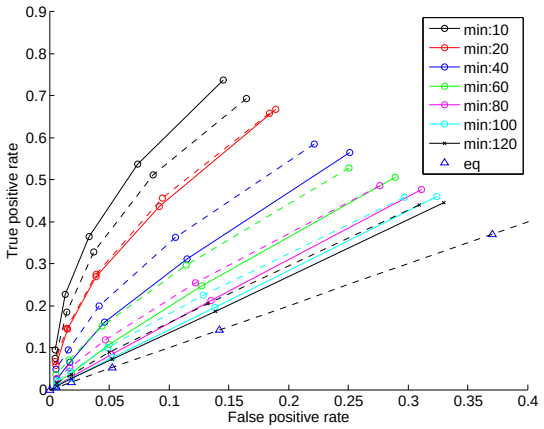

FIG. 4. AL prediction skill in terms of the ROC metric, i.e., TPR versus FPR, for 2003. Results are shown for the model trained with both AL and vB_z data (dashed) and the model trained with AL data alone (solid). In each curve, circles denote the thresholds indicated in the main text from the smaller values (with larger TPR) to the larger values (with smaller TPR). Triangles indicate the ROC values for a trivial random forecast lying on the TPR = FPR line.

.

of predictions for the AL index to lie below the threshold a. In Fig. 4, we show the true positive rate (TPR) versus false positive rate (FPR), which are defined as follows. The TPR is defined as the ratio between the number of times the prediction that the AL index below a is correct and the number of times the AL index is below a. The FPR is defined as the ratio between the number of times the prediction that the AL index is below a is incorrect and the number of times the AL index is above a. For a skillful prediction TPR should be close to 1 while FPR should be close to 0. From the training data, we can count that the AL index is below the thresholds $a = -153$, -319, -485, and -652 nT for 37.0%, 14.3%, 5.24%, and 1.81% of the time, respectively (see the triangles in Fig.4). For a prediction scheme based on drawing randomly from the training dataset, then we predict 37% of the time that the AL index is below the threshold $a = -153$ nT regardless of whether the index is below the threshold or not. For this trivial prediction scheme, the TPR and FPR of the AL index to be below that threshold are both equal to 0.37. With a similar strategy, the TPR and FPR of the AL index to be below the thresholds $a = -319$ and -485 nT are 0.143 and 0.0524, respectively. In fact, the ROC curve for this prediction strategy is given by the TPR = FPR diagonal.

For diffusion forecasting, the TPR is simply the average of $p_t^{\mathrm{AL}}(a)$ on timestamps where the AL index is below a in the verification period. Similarly, the FPR is the average of $p_t^{\mathrm{AL}}(a)$ on timestamps where the AL index is above a. Figure 4 shows the ROC curves for both the bivariate and univariate models. As expected, the TPR/FPR ratios decrease with increasing leads. An important observation is that the TPR values from our nonparametric models always exceed the corresponding FPRs, which implies that they are more skillful than randomly drawing from the training dataset. Note that for $t > 20$ minutes the bivariate model gives higher TPR (as well as lower FPR) than the univariate model. This long-term improvement occurs despite the fact that at short leads (e.g., 10 minutes) the univariate model performs better; this is due to apparently better initial conditions with the univariate filter. A similar overall behavior is also observed for the year 2000 dataset (see SI).

In this paper, we have studied AL index prediction with the diffusion forecast model—a nonparametric autonomous model that approximates the Fokker-Planck equation of the underlying processes at arbitrary lead times. The forecasting outputs are probabilities that the AL index is below a user-defined threshold, which indicates the occurrence of a substorm. We found that the model trained with both the AL index and SW data time series gives higher (lower) probabilities on time intervals when the AL index is lower (higher) than the given threshold. This high prediction skill becomes more apparent at lead times 40–80 minutes. We also assessed the forecast skill with the ROC metric and found that diffusion forecasting has significantly higher skill than trivial forecasts based on random draws from the training data set, even for longer lead times such as 60–120 minutes.

This research was supported by the ONR MURI grant N00014-12-1-0912. D.G. received support from ONR grant N00014-14-1-0150 and ONR YIP grant N00014-16-1-2649. M.G. received support from NSF grant AGS-1303646. J.H. received support from ONR grant N00014-16-1-2888 and NSF grant DMS-1317919. J.H. also thanks Kayo Ide who brought this problem to his attention.

SUPPLEMENTARY INFORMATION

DATA SOURCES

AL index data were retrieved from the Kyoto Geomagnetism Data Service (`http://wdc.kugi.kyoto-u.ac.jp/`, accessed on September 2016). Solar wind (v and B_z) data were provided by the ACE [8] spacecraft mission, and retrieved from the OMNIweb service (`http://omniweb.gsfc.nasa.gov/`, accessed on September 2016).

KERNEL ALGORITHM FOR GENERATING THE BASIS FUNCTIONS

Following [3, 5, 9], we build our data-driven basis $\{\phi_k\}$ using eigenfunctions associated with the variable-

bandwidth kernel $K_\epsilon : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}_+$ defined as

$$K_\epsilon(y, y') = \exp\left(-\frac{\|y - y'\|^2}{\epsilon\sigma^{-1/\hat{m}}(y)\sigma^{-1/\hat{m}}(y')}\right). \quad (2)$$

In (2), d is the dimension of the ambient data space (i.e., $d = 2q$ or $d = q$ for bivariate and univariate diffusion forecast models, where q is the number of delays), ϵ is a positive bandwidth parameter, $\sigma : \mathbb{R}^d \mapsto \mathbb{R}_+$ is a positive bandwidth function (to be described momentarily), and $\hat{m}$ is an estimate of the intrinsic dimension m of the state space manifold M. Note that because M is embedded in $\mathbb{R}^d$ by the smooth observation map H (see main text), it inherits a smooth Riemannian metric g. We also note that for now the bandwidth parameter ϵ is unspecified, but we will select it below using an automatic procedure which also returns the dimension estimate $\hat{m}$.

Following the theory developed in [5], we use as a bandwidth function σ in (2) is an estimate of the sampling density of the data relative to the Riemannian measure associated with the observation map. That is, σ is a function chosen such that for $y = H(x)$, $\sigma(y)$ is equal to the sampling density $\frac{d\mu(x)}{d\nu}$ up to an $O(\epsilon)$ error, where ν is the Riemannian measure associated with g. Here, we compute $\sigma(y)$ using kernel density estimation as in [3, 5, 9], namely

$$\sigma(y) = \frac{1}{N(2\pi\tilde{\epsilon}r^2(y))^{\tilde{m}/2}} \sum_{i=0}^{N-1} \tilde{K}_{\tilde{\epsilon}}(y, y_i),$$
$$\tilde{K}_{\tilde{\epsilon}}(y, y_i) = \exp\left(-\frac{\|y - y_i\|^2}{\tilde{\epsilon}r(y)r(y_i)}\right),$$

where N is the number of samples in the training data, $\tilde{\epsilon}$ is a bandwidth parameter, $\tilde{m}$ an estimate of m (note that $\tilde{\epsilon}$ and $\tilde{m}$ are not necessarily equal to ϵ and $\hat{m}$, respectively). Moreover, $r(y)$ is a bandwidth function given by

$$r^2(y) = \frac{1}{k_{\text{nn}} - 1} \sum_{j=2}^{k_{\text{nn}}} \|y - y_{I(y,j)}\|^2,$$

where $I(y, j)$ is the index of the j-th nearest neighbor of y in the training dataset $\{y_i\}$, and k_{nn} is an integer-valued parameter controlling the size of the neighborhood used for kernel density estimation. The results in this work were obtained using the value $k_{\text{nn}} = 8$ as in [3, 5, 9], although the properties of the basis are not too sensitive on that choice.

To select appropriate values for the bandwidth parameters ϵ and $\tilde{\epsilon}$, as well as to estimate the intrinsic dimension m, we compute the kernel sums $S(\epsilon) = \sum_{i,j=0}^{N-1} K_\epsilon(y_i, y_j)/N^2$ and $\tilde{S}(\tilde{\epsilon}) = \sum_{i,j=0}^{N-1} \tilde{K}_{\tilde{\epsilon}}(y_i, y_j)/N^2$ and select ϵ and $\tilde{\epsilon}$ as the maximizers of $R(\epsilon) = \frac{d \log S(\epsilon)}{d \log \epsilon}$ and $\tilde{R}(\tilde{\epsilon}) = \frac{d \log \tilde{S}(\tilde{\epsilon})}{d \log \tilde{\epsilon}}$, respectively. As discussed in [5], this criterion can be interpreted as maximizing the "resolution" associated with the kernels K_ϵ and $\tilde{K}_\epsilon$. Moreover, as $N \to \infty$, the maximum values $\hat{m} = \max_\epsilon R(\epsilon)$ and $\tilde{m} = \max_{\tilde{\epsilon}} \tilde{R}(\tilde{\epsilon})$ are both equal to the intrinsic dimension m. Numerically, we approximate $R(\epsilon)$ and $\tilde{R}(\tilde{\epsilon})$, using finite differences computed for equispaced values of $\log \epsilon$ and $\log \tilde{\epsilon}$, and set $\hat{m}$ and $\tilde{m}$ to the corresponding approximate maximum values. In our experiments, we obtain values of $\hat{m}$ and $\tilde{m}$ in the range 3.4–4.3.

Next, following [9], we normalize the kernel in (2) to construct a Markov operator $\hat{\mathcal{P}}_\epsilon$ acting on functions on M using the normalization procedure introduced in the diffusion maps algorithm [4]. This operator is represented by a discrete Markov kernel P_ϵ such that

$$P_\epsilon(x, x_i) = \frac{K'_\epsilon(x, x_i)}{\sum_{j=0}^{N-1} K'_\epsilon(x, x_j)},$$
$$K'_\epsilon(x, x_i) = \frac{K_\epsilon(y, y_i)}{\sum_{j=0}^{N-1} K_\epsilon(y, y_j)}, \quad (3)$$

where x is an arbitrary point in M, $y = H(x)$ is the corresponding data point in $\mathbb{R}^d$, and $y_i = H(x_i)$ is the training data point associated with the state $x_i \in M$. In particular, given a function $f \in L^2(M, \mu)$, we define $\hat{\mathcal{P}}_\epsilon f(x) = \sum_{j=0}^{N-1} P_\epsilon(x, x_i) f(x_j)$. By ergodicity, as $N \to \infty$, $\hat{\mathcal{P}}_\epsilon f(x)$ converges almost surely to $\mathcal{P}_\epsilon f(x)$, where $\mathcal{P}_\epsilon$ is a Markov operator with a smooth transition kernel relative to the invariant measure μ. Moreover, it can be shown that $\mathcal{P}_\epsilon f(x)$ admits the asymptotic expansion,

$$\mathcal{P}_\epsilon f(x) = f(x) - \epsilon\Delta f(x) + O(\epsilon^2),$$

uniformly on M, where Δ is the Laplace-Beltrami operator associated with the conformally transformed Riemannian metric $\tilde{g} = \omega g$ with $\omega = \left(\frac{d\mu}{d\nu}\right)^{2/m}$. Due to the conformal scaling by ω, the volume form $\tilde{\nu}$ of $\tilde{g}$ is equal to the invariant measure of the dynamics (i.e., $\frac{d\mu}{d\tilde{\nu}} = 1$). Thus, Δ is a self-adjoint operator on $L^2(M, \mu)$ and its eigenfunctions $\{\varphi_0, \varphi_1, \ldots\}$ provide an orthonormal basis of this space. An important property of these basis functions is that they are extrema of the Dirichlet energy functional

$$E(f) = \int_M \|\text{grad}_g f\|_g^2 \omega^{-1} \, d\mu, \quad (4)$$

where $\text{grad}_g : C^\infty(M) \mapsto TM$ is the gradient operator associated with the ambient space metric g, and $E(\varphi_k)$ is equal to the corresponding eigenvalue. Intuitively, we think of E as a measure of roughness of functions that (due to the ω^{-1} term in the integrand) penalizes large gradients with respect to g in regions of M with small sampling density—this property endows the φ_k with robustness against sampling errors. Moreover, ordering the φ_k in order of increasing Dirichlet energy, we interpret $\{\varphi_0, \ldots, \varphi_{\ell-1}\}$ as the least-rough ℓ-element orthonormal

set in $L^2(M,\mu)$ with respect to the Dirichlet functional in (4).

With these considerations in mind, we build our data-driven basis $\{\phi_0,\ldots,\phi_{\ell-1}\}$ using the leading ℓ eigenfunctions (ordered in order of increasing eigenvalue) of the operator $\hat{\mathcal{G}}_\epsilon f(x) = (f(x) - \hat{\mathcal{P}}_\epsilon f(x))/\epsilon$, which converges pointwise to $\Delta f(x)$ as $N\to\infty$ and $\epsilon\to 0$. This is equivalent to representing each ϕ_k by an N-dimensional column vector $\vec{\phi}_k = (\phi_k(x_0),\ldots,\phi_k(x_N))$ storing the eigenfunction values at the sampled points in M, and solving the matrix eigenvalue problem $P\vec{\phi}_k = \lambda_k\vec{\phi}_k$, where P is an $N\times N$ matrix with elements $P_{ij} = P_\epsilon(x_i,x_j)$ given by (3), and the eigenvalues λ_k are ordered in decreasing order. Numerically, we solve this problem by first computing the eigenvectors $\vec{\psi}_k$ of the symmetric matrix $S = D^{1/2}PD^{-1/2}$, where D is an $N\times N$ diagonal matrix with entries $D_{ii} = \sum_{j=0}^{N-1} K'_{ij}$, and then evaluating $\vec{\phi}_k = D^{-1/2}\vec{\psi}_k$. Moreover, for computational efficiency, we sparsify S by zeroing out the entries in each row which lie outside a pre-selected mutual affinity neighborhood of the corresponding data point. That is, we set S_{ij} to zero if $K_\epsilon(y_i,y_j)$ is not in the J largest values of $\{K_\epsilon(y_i,y_0),\ldots,K_\epsilon(y_i,y_{N-1})\}$ and $\{K_\epsilon(y_j,y_0),\ldots,K_\epsilon(y_i,y_{N-1})\}$. Due to the exponential decay of the kernel, this truncation does not impart an additional asymptotic error in the approximation of $\Delta f(x_i)$ by $\hat{\mathcal{G}}_\epsilon f(x_i)$ at fixed J.

Before proceeding, we note that [3, 5] perform a different kernel normalization procedure than the one described above and in [9]. In particular, they approximate a weighted Laplacian $\hat{\mathcal{L}}$ associated with the ambient space metric g, as opposed to the unweighted Laplace-Beltrami operator Δ associated with the conformally transformed metric $\tilde{g}$. The weighted Laplacian in [3, 5] is associated with the Dirichlet functional $\int_M \|\mathrm{grad}_g\, f\|_g^2\, d\mu$ which differs from (4) in that it does not feature the density-dependent term ω^{-1} in the integrand. However, both $\hat{\mathcal{L}}$ and Δ are self-adjoint on $L^2(M,\mu)$, and their eigenfunctions provide suitable orthonormal bases for diffusion forecasting.

AL PREDICTION RESULTS FOR 2003

Supporting information (SI) Figs. 5 and 6 show analogous prediction results for 2003 to Figs. 2 and 3 in the main text, but for the lower AL index threshold values $a=-485$ and -652 nT corresponding to the mean minus two and three standard deviations of the AL time series, respectively. As expected, the prediction probabilities for these thresholds are generally lower than the results for the $a=-319$ nT threshold shown in the main text, but they still exhibit significant increases when the AL verification signal crosses the chosen threshold. This suggests that the diffusion model has skill even for thresholds as low as three standard deviations, as is indeed confirmed in the ROC results in Fig. 4 in the main text. Moreover, comparing SI Fig. 6 to Fig. 3 in the main text indicates that the improvement of skill of the bivariate models over the univariate models is higher for the lower thresholds.

AL PREDICTION RESULTS FOR 2000

SI Figs. 7 and 8 show analogous prediction probability results to Fig. 2 in the main text and SI Fig. 5, respectively, but for the year 2000 dataset. The corresponding ROC results are shown in SI Fig. 9. As in the case of year 2003, the bivariate models trained with both AL and vB_z data for 2000 have generally higher prediction skill than the corresponding univariate models trained on AL index data alone, apart from small lead times ($t\lesssim 20$ minutes) where the univariate models appear to benefit from better initial conditions. Moreover, the improvement of skill of the bivariate models over the univariate ones is higher for the more challenging forecasts with the smaller (-485 and -651 nT) AL index thresholds.

AUXILIARY vB_z PREDICTION RESULTS

Since the diffusion forecast models constructed here are autonomous, the bivariate models trained with both the AL and vB_z data can also be used to predict vB_z. SI Fig. 10 shows the root mean squared (RMS) errors of vB_z prediction for years 2000 and 2003. Here, the RMS error is defined as

$$\mathbb{E}[z^{\mathrm{SW}}](t) = \int_M z^{\mathrm{SW}}\rho_t\, d\mu \approx \frac{1}{N}\sum_{i=0}^{N-1} z^{\mathrm{SW}}(x_i)\hat{\rho}_t(x_i),$$

where the average is taken over the 20000 sample verification period. Notice that the RMS errors relax to the corresponding steady-state values (i.e., the ensemble spread) beyond lead times of 200 minutes.

[1] D. Vasiliadis, A. J. Klimas, D. N. Baker, and D. A. Roberts, J. Geophys. Res. **100**, 3495 (1995); M. I. Sitnov, A. S. Sharma, and K. Papadopoulos, Phys. Rev. E **65**, 016116 (2001); A. Y. Ukhorskiy, M. I. Sitnov, A. S. Sharma, and K. Papadopoulos, J. Geophys. Res. **107**, 1369 (2002).

[2] V. V. Gavrishchaka and S. B. Ganguli, J. Geophys. Res. **106**, 29911 (2001).

[3] T. Berry, D. Giannakis, and J. Harlim, Phys. Rev. E. **91**, 032915 (2015).

[4] R. R. Coifman and S. Lafon, Appl. Comput. Harmon. Anal. **21**, 5 (2006).

[5] T. Berry and J. Harlim, Appl. Comput. Harmon. Anal. **40**, 68 (2016).

[6] T. Berry and J. Harlim, Phys. D **320**, 57 (2016).

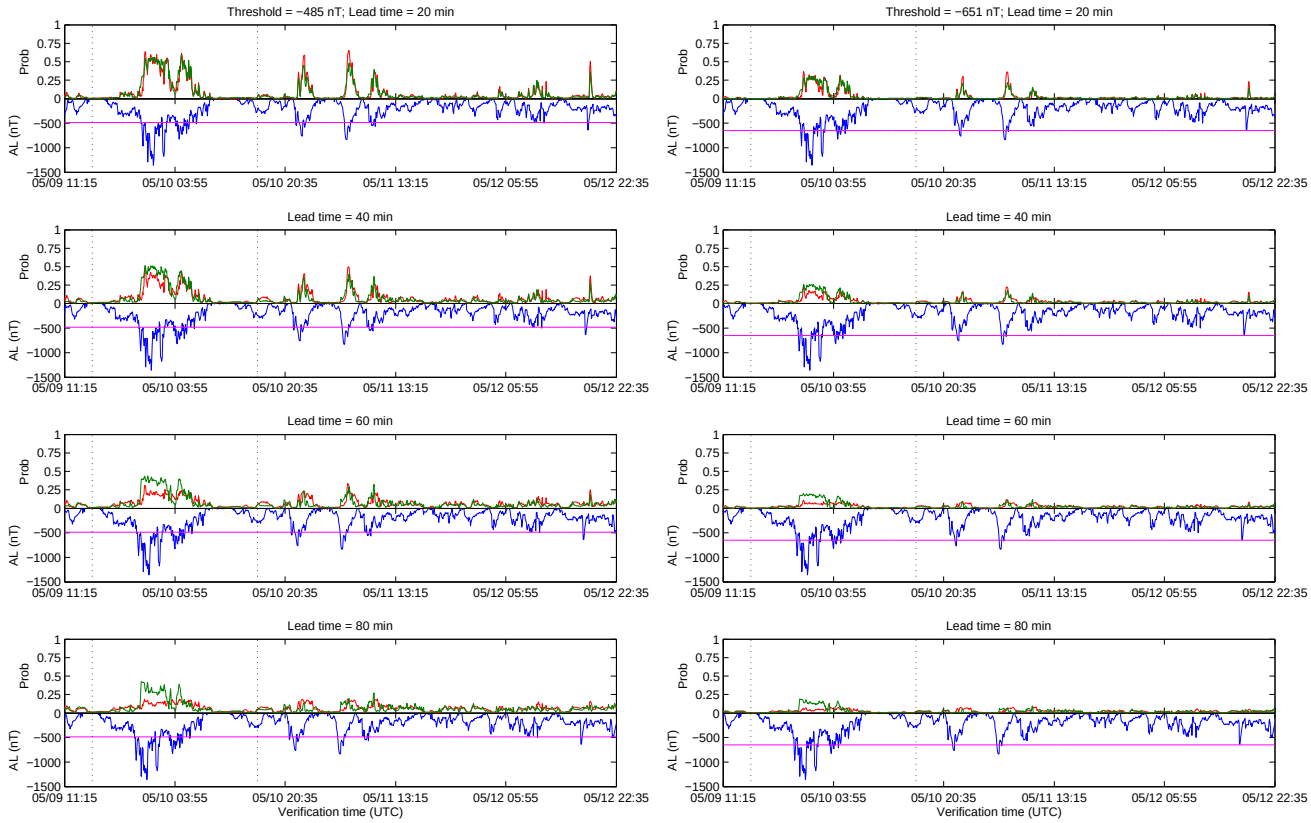

FIG. 5. Prediction of the AL time series for the interval 5/9/2003 11:15 UTC to 5/12/2003 22:35 UTC. The panels show the true time series (blue) and the prediction $p_t^{\mathrm{AL}}(a)$ for $a = -485$ nT (left) and -652 nT (right) at lead times $t = 20, 40, 60, 80$ minutes. Green and red lines show the predictions based on the bivariate and univariate models, respectively. Horizontal magenta lines show the AL index threshold a. The vertical doted lines indicate the interval shown in detail in SI Fig. 6.

[7] N. H. Packard *et al.*, Phys. Rev. Lett. **45**, 712 (1980); F. Takens, in *Dynamical Systems and Turbulence*, Lecture Notes in Mathematics, Vol. 898 (Springer, Berlin, 1981) pp. 366–381; T. Sauer, J. A. Yorke, and M. Casdagli, J. Stat. Phys. **65**, 579 (1991).

[8] E. C. Stone, A. M. Frandsen, M. R. A., E. R. Christian, D. Margolies, J. F. Ormes, and F. Snow, Space Sci. Rev. **86**, 1 (1998).

[9] D. Giannakis, "Data-driven spectral decomposition and forecasting of ergodic dynamical systems," (2016), arXiv:1507.02338.

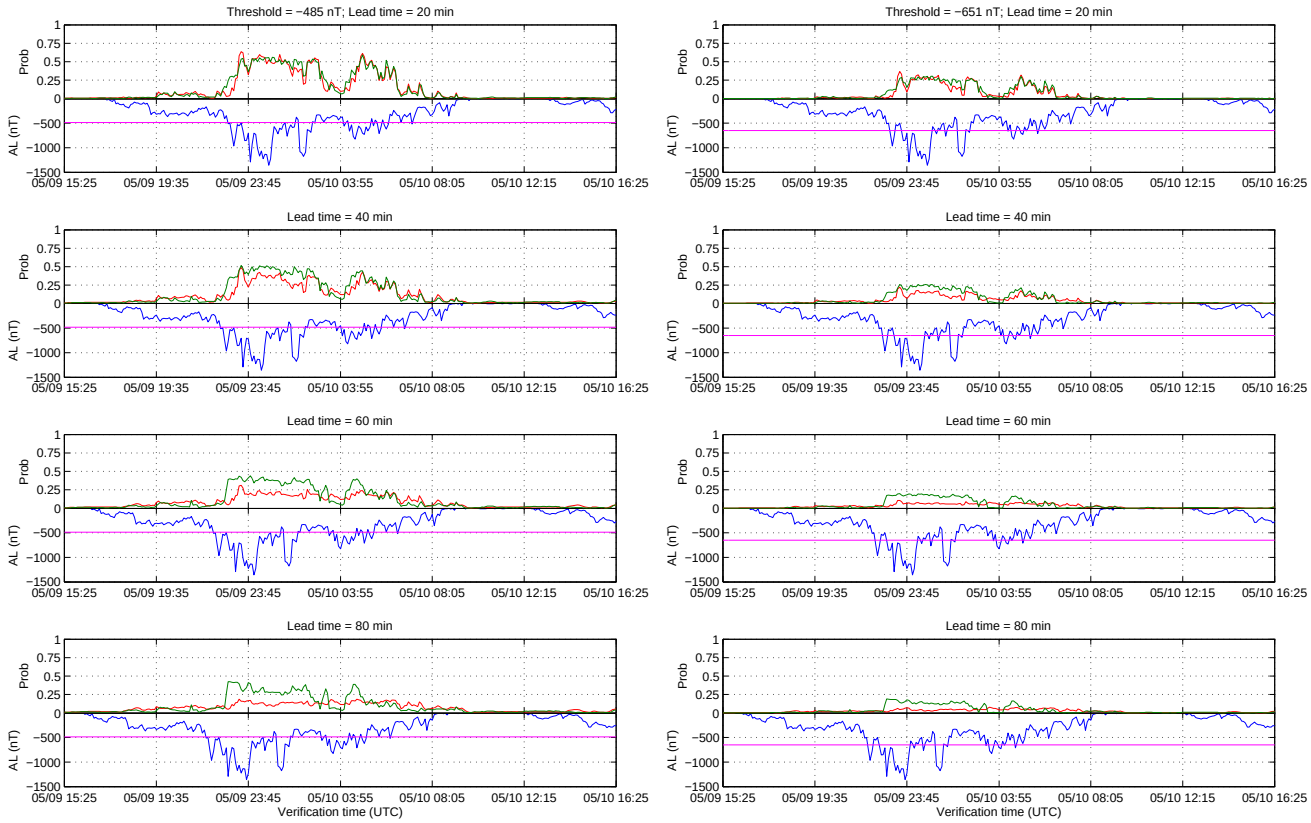

FIG. 6. Same as SI Fig. 5, but for the interval 5/9/2003 11:15 UTC to 5/12/2003 22:35 UTC.

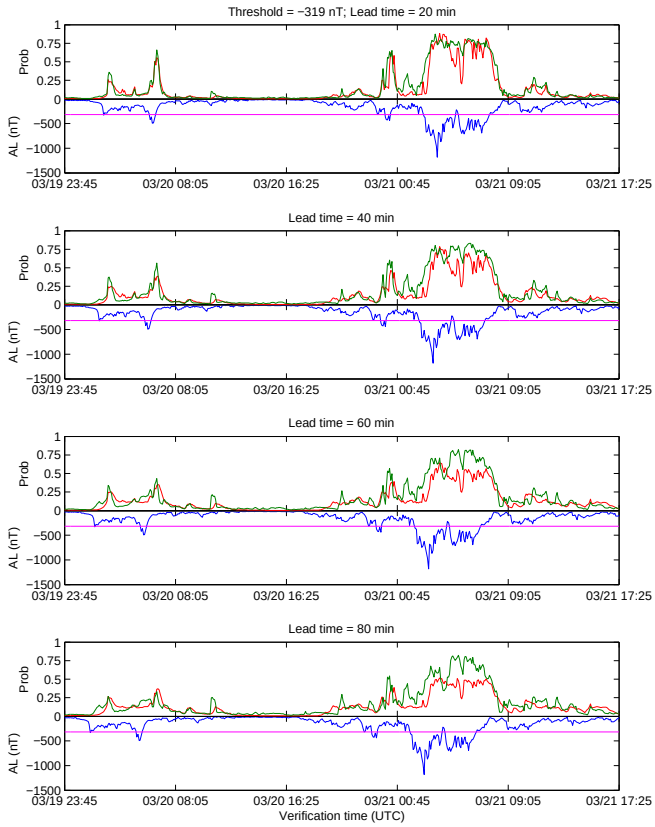

FIG. 7. Same as SI Fig. 5, but for the year 2000 dataset and the threshold $a = -319$ nT.

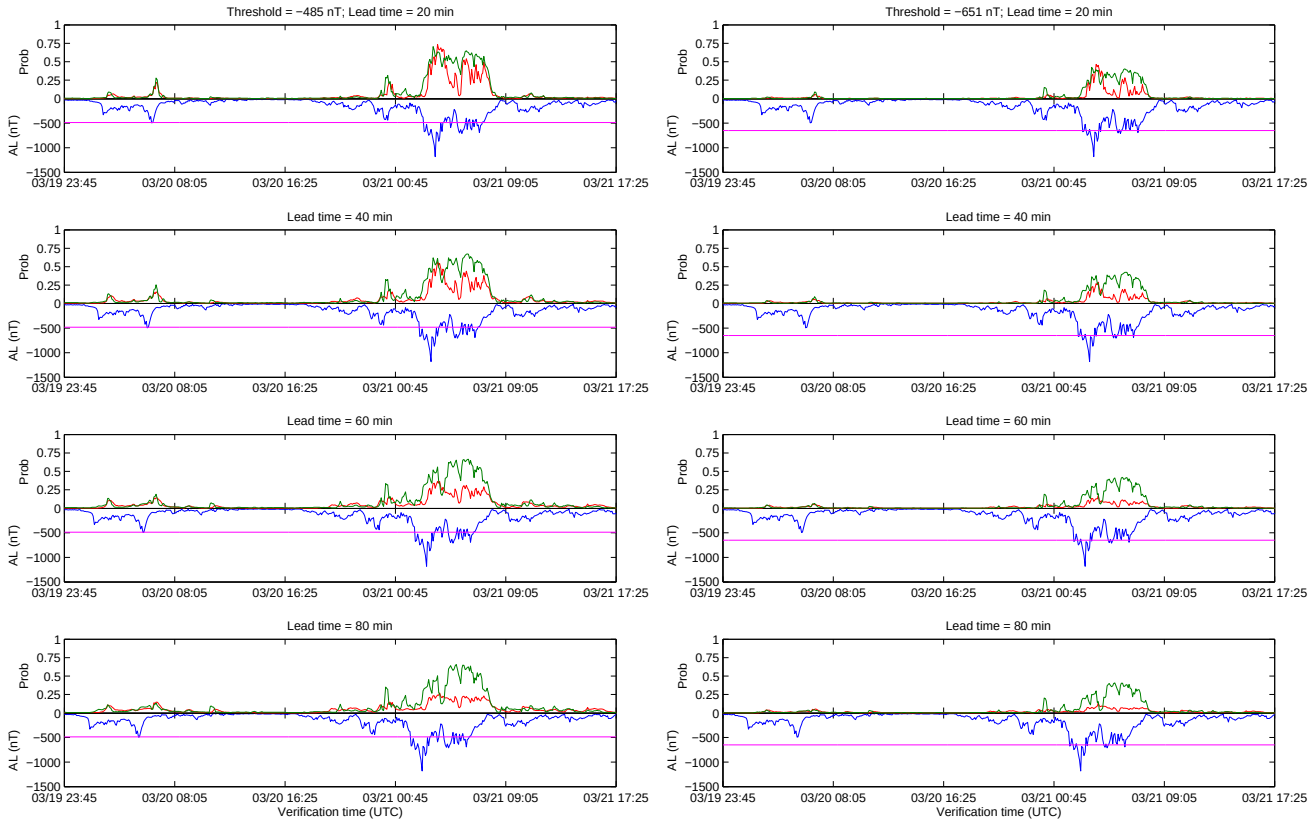

FIG. 8. Same as SI Fig. 8, but for the thresholds $a = -485$ nT (left) and -651 nT (right).

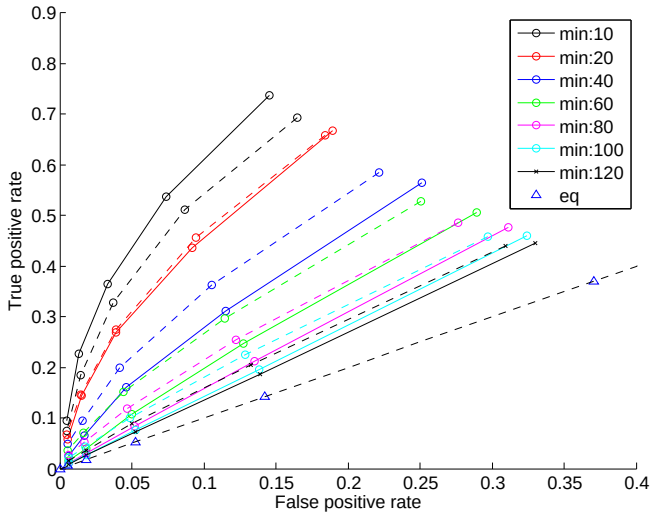

FIG. 9. AL prediction skill in terms of the ROC metric, i.e., TPR versus FPR, for 2000. Results are shown for the model trained with both AL and vB_z data (dashed) and the model trained with AL data alone (solid). In each curve, circles denote the thresholds indicated in the main text from the smaller values (with larger TPR) to the larger values (with smaller TPR). Triangles indicate the ROC values for a trivial random forecast lying on the TPR = FPR line.

.

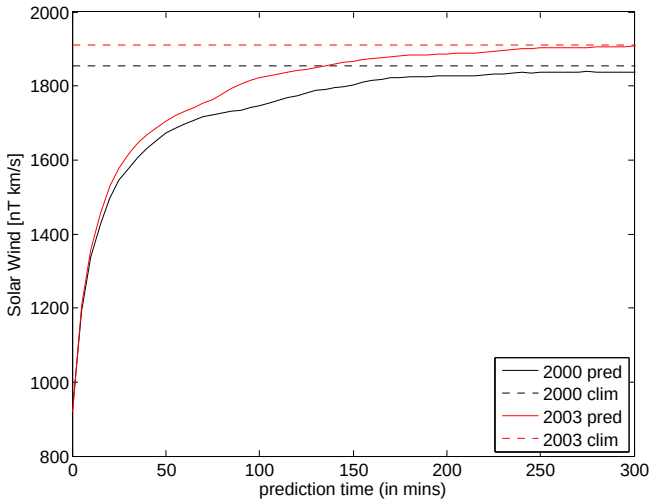

FIG. 10. vB_z prediction: RMS error as a function of lead time time.